arXiv:1202.5270v1 [quant-ph] 23 Feb 2012

Robust error bars for quantum tomography

Robin Blume-Kohout[1]

[1] *Theoretical Division (T-4/CNLS), MS-B258, Los Alamos National Laboratory, Los Alamos, NM 87545**

In quantum tomography[1], a quantum state or process is estimated from the results of measurements on many identically prepared systems. Tomography can never identify the state (ρ) or process ($\mathcal{E}$) exactly. Any point estimate is necessarily "wrong" – at best, it will be close to the true state. Making rigorous, reliable statements about the system requires *region estimates*. In this article, I present a procedure for assigning *likelihood ratio (LR) confidence regions*, an elegant and powerful generalization of error bars. In particular, LR regions are almost optimally powerful – i.e., they are as small as possible.

Quantum information processing relies on quantum hardware, including memory qubits and unitary or nearly-unitary quantum gates. These individual components must perform their allotted transformations with very high precision, especially for fault-tolerant quantum computing. The methods used to characterize and validate quantum devices are known, collectively, as *quantum tomography*. Tomography usually involves repeated independent measurements on N identically-prepared systems (referred to hereafter as "standard tomography"), but can also involve collective measurements on all N copies. Because state and process tomography are mathematically equivalent, this paper will focus on state tomography for the sake of clarity, with the understanding that all results can be extended straightforwardly to processes[13].

Tomography cannot identify ρ (the state produced by a quantum device) *exactly*, for precisely the same reason that flipping a coin N times cannot reveal its bias exactly. Any point estimate $\hat{\rho}$ has precisely zero probability of coinciding exactly with the true ρ, for there are infinitely many other states arbitrarily close to $\hat{\rho}$ and equally consistent with the data. To make a tomographic assertion about the device that *is* true – or at least true with high probability – we must report a *region* of states or processes (vis. Fig. 1).

Such regions are often constructed by attaching *error bars* to a point estimate. In quantum tomography, this approach suffers several drawbacks, some of which are illustrated in Fig. 2. Naïve error bars define an ellipsoidal shape (arbitrary), centered at the point estimate (suboptimal), which may include many unphysical states (inefficient). Worst of all, it is generally impossible to assign this ellipsoid any rigorous meaning – e.g., "The true state is within it, with probability at least 99%." The same problem applies to the other method used to date, *bootstrapping*[2] – which means generating a host of simulated datasets $\{D_k\}$ (either by resampling the real data, or by simulating measurements on a point estimate $\hat{\rho}$),

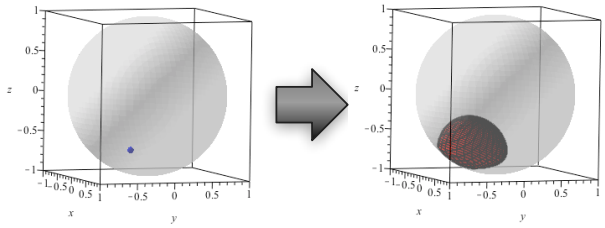

FIG. 1: *Point estimators*, like the maximum likelihood estimate $\hat{\rho}_{\mathrm{MLE}}$ shown on the left, cannot provide meaningful and rigorous statements about the true (but unknown) state ρ. But if we replace point estimators with *region estimators*, like the likelihood-ratio confidence region shown on the right, then the region $\hat{\mathcal{R}}$ defines an assertion – "ρ lies within $\hat{\mathcal{R}}$ with 90% certainty" – that is rigorously valid. The estimates shown here came from simulated measurements on 60 copies of a single-qubit state, with 20 measurements each of $\sigma_x, \sigma_y, \sigma_z$ yielding +/- counts of 7/13, 9/11, and 3/17 (respectively).

then reporting the variance of the corresponding point estimates $\{\hat{\rho}_k\}$. The underlying problem is that bootstrapping and naïve error bars both represent *standard errors* – the variance of a point estimator. Unfortunately, the point estimators used in quantum tomography are all *biased*, and standard errors for biased estimators do not reliably represent uncertainty[14] about the true ρ.

Happily, all of these issues can be resolved with a re-

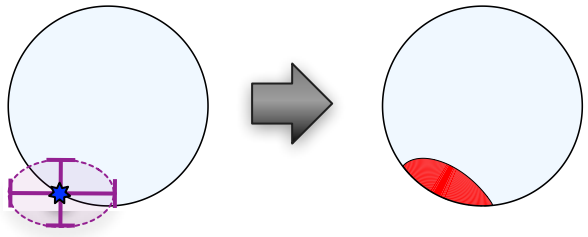

FIG. 2: General region estimates – adapted to the data, and constructed so as to minimize volume – can be far more powerful, useful, and reliable than traditional "error bars". As illustrated here, a valid confidence region need not be: (i) ellipsoidal or rectangular, (ii) centered at a point estimate, or (iii) aligned with the axes defined by whatever observables were measured. The figure on the right shows a cross-section of a 1-qubit LR confidence region, while on the left the smallest traditional error bars with the same coverage probability are shown. The LR region is noticeably smaller, and includes only valid states. Although in this case, the LR region could be reasonably approximated by the intersection between the error ellipsoid and the Bloch sphere, this is not always the case.

markably simple construction. *Likelihood ratio (LR) confidence regions* (see Fig. 3 for some examples) generalize the notion of error bars, providing data-adapted (not necessarily ellipsoidal) regions that:

1. contain the true state with guaranteed, high, and user-specified probability;

2. are, on average, smaller (and thus more powerful) than almost any other construction; and

3. are simple to define and construct.

Definition 1. *Given observed data D, the **likelihood** is a function on states given by $\mathcal{L}(\rho) = Pr(D|\rho)$. The **log likelihood ratio** is a function on states given by $\lambda(\rho) = -2\log\left[\mathcal{L}(\rho)/\max_{\rho'}\mathcal{L}(\rho')\right]$. Given data D, the **likelihood ratio region** with confidence α is $\hat{\mathcal{R}}_\alpha(D) = \{\text{all } \rho \text{ such that } \lambda(\rho) < \lambda_\alpha\}$, where λ_α is a constant (see below) that depends on the desired confidence α and the Hilbert space dimension d.*

It should be obvious that the threshold value λ_α plays a critical role in this construction. Increasing λ_α increases the size of the LR region $\hat{\mathcal{R}}_\alpha$, which in turn increases the estimator's *coverage probability* – the probability that $\hat{\mathcal{R}}(D)$ will contain ρ – but reduces its *power* (since large regions imply less about ρ). So λ_α should be set to the smallest value that ensures coverage probability at least α. This optimal value depends on α and the size of the system we are measuring. It is hard to compute exactly, but two upper bounds are provided in Eq. 14 and Lemma 1. Either will guarantee coverage probability at least α, at the cost of slightly increasing the regions' size.

The remainder of this paper attempts to answer three natural questions, in order of increasing technicality. First, "What is a confidence region, and how does it generalize 'error bars'?" Next, "Why is the LR construction an especially good one?" Finally, "How do we choose the threshold?" The concluding discussion section addresses a few other questions, especially the relationship to related work by Christandl and Renner[3], and how to use and describe LR regions.

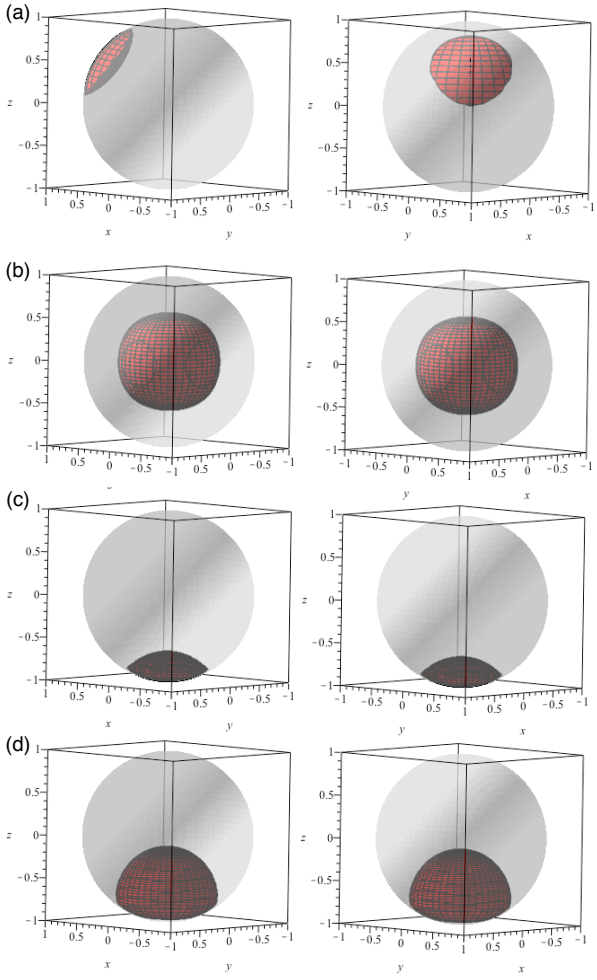

FIG. 3: Four examples of the shapes that LR regions can take. Each example is a 90% confidence region for a qubit, based on 60 measurements divided equally among the three Pauli observables. Rows (a-d) correspond to four distinct datasets. Left and right columns are different views of the same region.

I. REGION ESTIMATORS AND CONFIDENCE REGIONS

Error bars around a point estimate define a region estimate, but region estimates do not need to be associated with a point estimate. After seeing the data (D), we can assign a region $\hat{\mathcal{R}}(D)$ of whatever shape and size is necessary to achieve our goals. So what are these goals?

First and foremost is *coverage probability*. By assigning a region, we assert that the unknown ρ is within it. This had better be true with very high probability. Coverage probability (α) is the probability that $\hat{\mathcal{R}}(D)$ does indeed contain ρ. It would be very satisfying indeed if we could assert "Given the observed data, the probability that ρ is in $\hat{\mathcal{R}}(D)$ is α," i.e.

$$Pr(\rho \in \hat{\mathcal{R}}(D)|D) \geq \alpha.$$

Unfortunately, this assertion[15] *requires* assigning a *prior* probability to "$\rho \in \hat{\mathcal{R}}(D)$". Different agents (e.g., a scientist, a skeptical reader, and a funding agency) will generally disagree about this prior, and therefore about $Pr(\rho \in \hat{\mathcal{R}}(D)|D)$.

So, instead, we make a subtly different assertion: "The region $\hat{\mathcal{R}}(D)$ that we assign will contain ρ with probability at least α," i.e.

$$Pr(\rho \in \hat{\mathcal{R}}(D)) \geq \alpha. \quad (1)$$

This assertion is made *before* the data D are observed. Once we know D, $\hat{\mathcal{R}}(D)$ is fixed, and the most that we can say is "This region was obtained by a procedure that 'works' almost always (i.e., with probability $\geq \alpha$)". Now

(after the estimate), α is not the probability of success – for success is no longer a random variable. Instead, α quantifies our *confidence* in the estimate, and so an estimation procedure satisfying this condition is a *confidence region* estimator.

Confidence is a property of the entire estimator – the map from data to regions – rather than of a particular region estimate. It depends on the unknown state – but, *mirabile dictu*, we can place relatively tight prior-*independent* bounds on it,

$$Pr(\rho \in \hat{\mathcal{R}}(D)) \;\geq\; \min_{\rho} Pr(\rho \in \hat{\mathcal{R}}(D)|\rho).$$

A confidence region estimator with confidence α satisfies

$$\Pr(\rho \in \hat{\mathcal{R}}(D)|\rho) \geq \alpha \;\forall\; \rho.$$

It's important to understand that confidence regions do not provide probabilistic statements about any single run of the experiment. Once the data are taken, the estimator either succeeded or failed, and there is no way to assign a probability to its success without choosing a prior. In any given experiment, what can be said is "We applied a technique which is guaranteed to yield a region containing the true ρ at least α of the time – no matter what the unknown ρ is."

II. OPTIMALITY

Confidence regions are a basic statistical construct, especially for scalar parameters (where they are known as confidence *intervals*). Yet even for 1-dimensional intervals, there exist many distinct constructions, and no consensus on the "best" choice. Note that designing a high-confidence region estimator isn't hard. For example, the estimator $\hat{\mathcal{R}}(D) = \{\text{all states}\}$ has coverage probability $\alpha = 1$. The challenge is to design one that is *powerful* – i.e., assigns small regions (which correspond to powerful hypotheses, because they rule out many states).

The likelihood ratio construction in Definition 1 was introduced in 1989 in the context of particle physics by Feldman and Cousins[4], and appears to have been part of statistical folklore before that. What has *not* appeared to date is a compelling argument why LR regions are particularly good. In this section, I provide such a justification, comprising: (1) a proof that the most powerful confidence region estimators are *probability ratio* (PR) estimators (a similar result was proven by Evans et al[5]; the treatment here is self-contained), and (2) a heuristic argument that, among the family of PR estimators, LR estimators have nearly-optimal worst-case behavior.

First, we need to quantify the power of any given region estimator $\hat{\mathcal{R}}(\cdot)$. Smaller regions are clearly more powerful, and in this work I will quantify a region's power by its *volume*,

$$V(\hat{\mathcal{R}}) = \int_{\rho \in \hat{\mathcal{R}}} \mathrm{d}\rho. \tag{2}$$

Choosing a particular volume measure $\mathrm{d}\rho$ could be controversial. Remarkably, the construction in Definition 3 is optimal for *any* measure $\mathrm{d}\rho$!

Now, the volume of the assigned region is itself a random variable, depending on the data. Averaging over datasets yields an expected volume,

$$\overline{V}(\rho) = \sum_D Pr(D|\rho) V(\hat{\mathcal{R}}(D)),$$

which is a function of ρ. Since ρ is (by definition) unknown, we can quantify the estimator's performance either by *worst-case* (maximum) or *average* volume. To average, we must choose a measure $\mu = P(\rho)\mathrm{d}\rho$ over ρ (which need *not* be related in any way to $\mathrm{d}\rho$). Optimal average performance is achieved by that measure's *Bayes estimator*[6].

Definition 2. *Given a cost function V, the **Bayes estimator** for a given measure μ (over states) is the estimator with the smallest average expected cost w/r.t. μ.*

But the Bayes estimator for one measure might have very bad performance for another measure. So if we are not sure what measure to choose, we have an alternative: choose an estimator that minimizes worst-case performance, $\max_\rho \overline{V}(\rho)$. This defines the *minimax* estimator[6]. Happily (and perhaps surprisingly), these two approaches are intimately related by a basic theorem of decision theory:

Theorem 1. ***(Minimax-Bayes duality[6,7])*** *The minimax estimator (for a given cost function) is, under mild regularity conditions, also the Bayes estimator for some measure μ, known as the* least favorable prior.

So to find the minimax estimator – which is appealing because it has the best possible *guaranteed* performance – we will focus first on optimizing average performance

$$\left\langle \overline{V} \right\rangle_P = \int \overline{V}(\rho) P(\rho) \mathrm{d}\rho,$$

for an arbitrary measure $P(\rho)\mathrm{d}\rho$. (N.B. In spite of this notation, the averaging measure $P(\rho)\mathrm{d}\rho$ and volume measure $\mathrm{d}\rho$ are *completely* independent! For example, Eq. 4 doesn't depend on the choice of volume measure $\mathrm{d}\rho$.)

A region estimator $\hat{\mathcal{R}}(\cdot)$ can be represented by a connection relation: for each dataset D and state ρ, we say that ρ is "connected" to D ($\rho \sim D$) iff $\rho \in \hat{\mathcal{R}}(D)$. The average expected volume is then given by

$$\begin{aligned} \left\langle \overline{V} \right\rangle &= \int \overline{V}(\rho) P(\rho) \mathrm{d}\rho \\ &= \int \sum_D Pr(D|\rho) V(\hat{\mathcal{R}}(D)) P(\rho) \mathrm{d}\rho \\ &= \sum_D Pr(D) V(\hat{\mathcal{R}}(D)), \end{aligned} \tag{3}$$

where

$$Pr(D) \equiv \int Pr(D|\rho) P(\rho) \mathrm{d}\rho. \tag{4}$$

Now, we recall Eq. 2 for the volume of a region, and observe that the integral over $\rho \in \hat{\mathcal{R}}(D)$ is equivalent to an integral over $\rho \sim D$, which gives

$$\begin{aligned} \langle \overline{V} \rangle &= \sum_D Pr(D) \int_{\rho \sim D} \mathrm{d}\varphi \\ &= \int_{\text{all } \rho} \left(\sum_{D \sim \rho} Pr(D) \right) \mathrm{d}\rho. \end{aligned} \tag{5}$$

Now, recall that we want to minimize $\langle \overline{V} \rangle$ (Eq. 5) subject to the constraint that

$$\sum_{D \sim \rho} Pr(D|\rho) \geq \alpha \text{ for each } \rho. \tag{6}$$

We can minimize each of the terms in the integral (Eq. 5) independently – because they are not coupled by the constraint. To do so, consider a simple cost/benefit analysis. Eq. 6 says that each state ρ must be connected to datasets whose total *conditional* probability $Pr(D|\rho)$ is at least α. But Eq. 5 says that each such connection comes at a cost, given by the *unconditional* probability $Pr(D)$. To achieve total conditional probability α at minimum cost, we connect ρ to datasets in descending order of benefit/cost ratio, given by the *probability ratio* statistic

$$r(D;\rho) = \frac{Pr(D|\rho)}{Pr(D)}, \tag{7}$$

down to a threshold $r_\alpha(\rho)$ that satisfies

$$\sum_{\text{all } D \text{ s.t. } r(D;\rho) > r_\alpha(\rho)} Pr(D|\rho) \geq \alpha. \tag{8}$$

Inverting this relationship (to define which states are connected to a given D, rather than the other way around) yields probability ratio (PR) region estimators:

Definition 3. *The PR region estimator for an averaging measure $\mu = P(\rho)\mathrm{d}\rho$, with confidence level α, is defined by $\hat{\mathcal{R}}(D) = \{\text{all } \rho \text{ such that } Pr(D|\rho)/Pr(D) \geq r_\alpha(\rho)\}$, with $Pr(D)$ and $r_\alpha(\rho)$ given by Eqs. 4,8.*

This prescription is an exact solution to the problem of minimum-average-volume confidence regions – and, as advertised, it does not depend on what measure $\mathrm{d}\rho$ is used to defined volume!

PR estimators are interesting in themselves. In particular, the estimator that unconditionally "works best" for $\rho = \rho_0$ (for any given ρ_0) is the PR estimator for the measure $\mu = \delta(\rho - \rho_0)$. An especially confident experimentalist who believes that her apparatus really is producing ρ_0 might choose this estimator. Despite the radical "prior", this estimator still assigns valid confidence regions, on which a skeptical third party can rely. The experimentalist's extreme confidence is reflected only in this manner: the datasets D that typically occur when $\rho = \rho_0$ yield relatively small regions, while datasets D that are improbable given $\rho = \rho_0$ (but might appear with high probability for other states) yield enormous regions. Thus, if ρ really is ρ_0, then the experimentalist is rewarded with (moderately) small regions... but if she is wrong and ρ is very different from ρ_0, then the assigned region will probably be so large as to imply virtually nothing about ρ.

So while the extreme choice $\mu = \delta(\rho - \rho_0)$ is a valid one, it's unwise in practice. It *does* play an important role by establishing an absolute (and tight) lower bound on $\overline{V}(\rho_0)$. But in practice, any sane experimentalist or analyst will choose a more balanced estimator – one that performs well even in the worst case. This (as noted above) is the minimax estimator, and is the PR estimator for the least favorable prior.

Finding exact LFPs is arduous and tricky at best[5]. Worse yet, the LFP will depend (perhaps sensitively) on the exact volume measure $\mathrm{d}\rho$. In order to circumvent this task (which remains a good challenge for future research), let's apply a simple heuristic ansatz to choose μ.

Suppose that we choose some μ, and then ρ is chosen by an adversary so as to maximize $\overline{V}(\rho)$. To do so, the adversary would look for a dataset D_0 and a state ρ such that $Pr(D_0|\rho) \gg Pr(D_0)$ (the latter is determined by $P(\rho)\mathrm{d}\rho$). If $Pr(D_0)$ is relatively small, then its "cost" will be relatively low, and many states ρ' will be connected to it – which means that $\hat{\mathcal{R}}(D_0)$ will be large. But because $Pr(D_0|\rho)$ is relatively large, D_0 will occur relatively often, and $\overline{V}(\rho)$ will be large.

Avoiding this vulnerability is simple: ensure that $Pr(D|\rho)/Pr(D)$ is not too large for any ρ. This means choosing μ so that

$$Pr(D) \propto \max_\rho Pr(D|\rho). \tag{9}$$

With this choice, the probability ratio statistic (Eq. 7) becomes the *likelihood ratio* statistic,

$$r = \frac{Pr(D|\rho)}{Pr(D)} \to \frac{Pr(D|\rho)}{\max_{\rho'} Pr(D|\rho')} = \frac{\mathcal{L}(\rho)}{\mathcal{L}_{\max}(\rho)} \equiv \Lambda, \tag{10}$$

which defines likelihood-ratio (LR) regions (Definition 1) as a special case of PR regions. We use $\lambda = -2\log\Lambda$, rather than Λ, to maintain a convenient connection with [extensive] previous work on likelihood ratios.

III. THE THRESHOLD

For a generic PR region estimator, the threshold value of the statistic (r_α) depends significantly on ρ. We could define LR regions in the same way, using a ρ-dependent threshold $\lambda_\alpha(\rho)$. But one of the special attributes of the LR statistic is that, unlike generic PR statistics, its distribution is approximately independent of ρ (as shown in Fig. 4). Proving this independence exactly depends on a Gaussian approximation that does not quite hold even

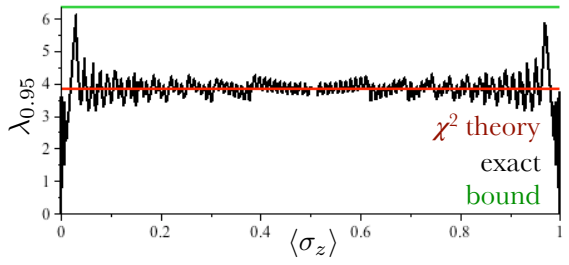

FIG. 4: This figure shows the tight state-*dependent* cutoff $\lambda_\alpha(\rho)$ for a particularly simple case: a qubit measured only in the σ_z basis (or, alternatively, a classical coin). The cutoff depends only on $\langle\sigma_z\rangle$, so it can be plotted easily and compared to (i) the χ^2 value, around which it fluctuates, and (ii) the upper bound given by Eq. 14. While using the state-dependent cutoff would yield (slightly) smaller regions, there is a concomitant loss of simplicity, elegance, and convenience (because the regions become nonconvex).

as $N \to \infty$, but it does hold approximately. So while *exactly* optimal regions require calculating a ρ-dependent threshold numerically (see Fig. 4), we can replace $\lambda_\alpha(\rho)$ with a constant lower bound λ_α satisfying

$$\lambda_\alpha \geq \lambda_\alpha(\rho) \text{ for all } \rho \tag{11}$$

and obtain a simpler and more elegant estimator that is still rigorously correct, and only sacrifices a little power. This simplifies matters – but we still need to set λ_α!

Coverage probability and region size both increase with λ_α. So we want to set it as low as possible (to ensure powerful regions) while maintaining coverage probability at least α. $\hat{\mathcal{R}}_\alpha$ will include ρ if and only if $\lambda(\rho) < \lambda_\alpha$. In principle, we can compute the distribution of $\lambda(\rho)$ for every possible ρ, define a *complementary cumulative distribution function* (CCDF)

$$F(\lambda_\alpha|\rho) = Pr(\lambda(\rho) > \lambda_\alpha|\rho),$$

and then solve the equation $\max_\rho F(\lambda_\alpha|\rho) = 1-\alpha$ for λ_α. In practice, computing F is hard, so instead we use upper bounds on F (as shown in Fig. 5) to set λ_α, ensuring coverage probability $\geq \alpha$ at a small cost in power. Two valid and useful (though not tight) bounds are given in Eq. 14 (whose derivation is rather arduous, and will be published elsewhere) and Lemma 1 (proven herein).

If the data had a Gaussian distribution, then $\lambda(\rho)$ would be a χ^2_k random variable, where k (the number of degrees of freedom) is the number of linearly independent observables measured, and is equal to d^2-1 for an informationally complete set of measurements. The corresponding CCDF is independent of the true state ρ, and given in terms of an upper incomplete Gamma function

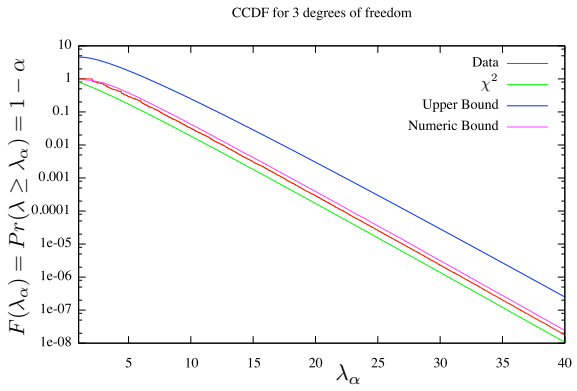

FIG. 5: Three different bounds on the complementary cumulative distribution function or CCDF $[F(\lambda_\alpha)]$ for the log-likelihood ratio derived from qubit state tomography (wherein ρ has $K = 3$ degrees of freedom). The horizontal axis is the cutoff value λ_α, while the vertical (log-scale) axis is the associated failure probability $1-\alpha = F(\lambda_\alpha)$. "**Data**" correspond to an exhaustive numerical calculation of the CCDF, and are compared with the χ^2 approximation (Eq. 12) and the upper bound given in Eq. 14. The plot labeled "**Numeric bound**" is a hybrid calculation where method used to derive Eq. 14 is augmented by calculating one hard-to-approximate quantity numerically; its excellent agreement with data supports Eq. 14, and suggests that it can be improved.

by

$$F_{\chi^2_k}(\lambda_\alpha) = \frac{\gamma(k/2, \lambda_\alpha/2)}{\Gamma(k/2)} \tag{12}$$

$$\approx \frac{1}{(\frac{k}{2}-1)!}\left(\frac{\lambda_\alpha}{2}\right)^{k/2-1} e^{-\lambda_\alpha/2}, \tag{13}$$

where the second line is valid as $\lambda_\alpha \to \infty$. Unfortunately, tomographic data are multinomial, not Gaussian, and this ansatz is too optimistic. Using it yields a coverage probability that is α only on average, and can be much lower for some ρ. A much more arduous calculation (to be published separately) yields an upper bound

$$F(\lambda_\alpha) \leq F_{\chi^2_k} + e^{-\lambda_\alpha/2}\left[\left(1+\frac{\sqrt{3e\lambda_\alpha}}{\pi}\right)^k - \frac{\sqrt{3e\lambda_\alpha}}{\pi}^k\right]$$

$$\to F_{\chi^2_k} + e^{-\lambda_\alpha/2} k\sqrt{\lambda_\alpha}^{k-1} \text{ when } \lambda_\alpha \to \infty, \tag{14}$$

that is valid whenever the data D are obtained from independent measurements on identically prepared copies of ρ (the standard tomographic setup). Figure 5 compares these bounds with an exhaustive numerical calculation of $F(\lambda_\alpha)$ for the $k = 3$ degrees of freedom found in single-qubit tomography.

A simpler (but looser) bound that applies to *any* data, including joint measurements[16], is

Lemma 1. *For any measurement on N copies of a d-dimensional system, $F(\lambda_\alpha) \leq N^{d^2-1}e^{-\lambda_\alpha/2}$.*

Proof. A POVM measurement $\mathcal{M} = \{E_k\}$ (where $E_k \geq 0$ and $\sum_k E_k = \mathbb{1}$) is performed on the state $\rho^{\otimes N}$, on the N-copy Hilbert space $\mathcal{H}^{\otimes N}$. By Schur's Lemma, we may decompose the Hilbert space as $\mathcal{H}^{\otimes N} = \bigoplus_\nu \mathcal{U}_\nu \otimes \mathcal{P}_\nu$, where $\mathcal{U}_\nu$ and $\mathcal{P}_\nu$ are irreducible representation spaces of the unitary group $SU(d)$ and the permutation group S_N (respectively). Because the state is permutation-symmetric, it is maximally mixed on the $\mathcal{P}_\nu$ factors, and therefore this measurement is equivalent to a measurement $\mathcal{M}' = \{E'_k\}$ on the much smaller Hilbert space $\bigoplus_\nu \mathcal{U}_\nu$, whose dimension is at most $M = N^{d^2-1}$.

Next, let the probability of event k given ρ be p_k, and the *maximum* probability of event k be q_k, so that $\lambda = -2\log(p_k/q_k)$. Rewriting this gives $p_k = q_k e^{\lambda/2}$. Now, $q_k \leq Tr(E'_k)$, and $\sum_k E'_k = \mathbb{1}$, so $\sum_k q_k \leq M$. The probability that $\lambda \geq \lambda_\alpha$ is $Pr(\lambda \geq \lambda_\alpha) = \sum_{\text{all } k \text{ s.t. } \lambda \geq \lambda_\alpha} p_k$. Each term is $p_k = q_k e^{-\lambda/2} \leq q_k e^{-\lambda_\alpha/2}$, so the sum is upper bounded by $M e^{-\lambda_\alpha/2} = N^{d^2-1} e^{-\lambda_\alpha/2}$. □

Equation 14 does not grow with the number of samples measured (N), so it will generally be much tighter than Lemma 1 – but is harder to derive. And in many cases Lemma 1 may be fine; because $F(\lambda_\alpha)$ decreases exponentially in all cases, the N^{d^2-1} factor will enlarge the region's size by at most polylog(N).

Since these bounds are loose, their use will produce confidence regions that are somewhat larger than necessary. We will want to know *how* excessively large they are! Fortunately, such a tool is ready to hand. The χ^2 approximation to $F(\lambda_\alpha)$, given in Eq. 12, gives a *lower* bound[17] on λ_α. This suggests a simple test:

1. Define a confidence region $\hat{\mathcal{R}}(D)$ using a value of λ_α obtained by solving Eq. 14 or the equation in Lemma 1.

2. Define an "inner bound" region $\hat{\mathcal{R}}_{\min}(D)$ using λ_α obtained by solving Eq. 12.

3. If $\hat{\mathcal{R}}(D)$ and $\hat{\mathcal{R}}_{\min}(D)$ are relatively similar (i.e., $\hat{\mathcal{R}}(D)$ is not much bigger, and they lead to similar conclusions about the experiment) then there's nothing to be gained from using a tighter bound.

IV. DISCUSSION

Likelihood-ratio confidence regions define "error bars" for quantum tomography that are:

1. Rigorously guaranteed to capture the true state (or process) with controllable probability α,

2. Approximately as small as can be achieved (within that constraint),

3. Natural, convenient, and intuitive.

The third point summarizes several particularly nice properties of LR regions. They are convex for standard tomographic data (because the log likelihood is convex), so they can be manipulated and characterized using convex programming. Determining whether a given state ρ lies in $\hat{\mathcal{R}}(D)$ is even easier – just compare $\mathcal{L}(\rho)$ to $\mathcal{L}_{\max}$. And LR regions are, in a sense, a natural generalization of the popular maximum-likelihood (ML) point estimator. This view is actually backward; whereas ML point estimators have no finite-sample optimality properties, and may yield pathological results for quantum tomography, LR regions form a provably near-optimal region estimator. So a more accurate view is that the near-optimality of LR regions explains why ML point estimators often work well: the true state is usually in $\hat{\mathcal{R}}_\alpha(D)$, and $\hat{\mathcal{R}}_\alpha(D)$ is usually a neighborhood of $\hat{\rho}_{\text{MLE}}$.

Unlike the ellipsoidal regions implied by error bars, or the spherical ones implied by large deviation bounds[8,9], LR regions have shapes that are variable and data-adapted (see Fig. 3). This is a virtue – they can be much smaller than the best region of a fixed shape. But it can also be inconvenient. Many questions about ρ can be answered directly using the simple implicit description of $\hat{\mathcal{R}}(D)$ and convex programming (e.g. "Is ρ definitely separable?" or "What values of $\langle X \rangle$ can be ruled out?"). Sometimes, though, an explicit description of $\hat{\mathcal{R}}(D)$ is required. One can be produced with reasonable efficiency by sampling from the surface and calculating a minimum-volume bounding ellipsoid[10] or hypersphere[11]. This trades power (the approximated region is larger) for convenience.

Most of the error bars or regions used to date in quantum tomography are based on standard errors – i.e., the variance of a point estimator, usually the maximum likelihood estimator. As discussed in the introduction, such regions are not reliable (they may work in many cases, but not all!). However, a rigorously reliable solution was proposed[3] quite recently by Christandl and Renner. This excellent result deserves some discussion here. Although it was obtained entirely independently of the current work[18], it addresses a very similar problem and arrives at a solution that (in some ways) is closely related... via remarkably different methods!

Their main result is that confidence regions with confidence α can be constructed by (i) constructing a Bayesian *credible* region (a region containing $\alpha' \approx 1$ of the posterior probability) for a particular prior (Hilbert-Schmidt measure over quantum states), then (ii) enlarging the region slightly in a particular way. This proof elegantly synthesizes Bayesian and frequentist notions, by constructing regions that are simultaneously confidence regions *and* credible regions (at least for a particular prior). It suggests that the two approaches – while philosophically orthogonal – are deeply related in some fashion.

On the other hand, the method in Ref.[3] is not explicit. That is, it does not suggest *which* high-posterior-probability region to report. The natural choice is to choose the smallest such region, but it is not immediately obvious how to identify it. Moreoever, there is no obvious way to determine how powerful this procedure is

– i.e., whether it assigns regions that are, in some sense, smaller than those assigned by most or all other estimators.

These are exactly the strengths of the likelihood-ratio method given here. Definition 1 defines a specific, simple, and straightforward protocol, and we can easily analyze the expected volume of LR regions and show that they are optimal in a strong sense. The resulting regions are also relatively easy to characterize using known properties of the likelihood function (e.g., for standard tomographic data it is convex, and so are the LR regions). On balance, the LR method given here seems more practically useful at the present time. However, it's worth noting that the posterior distribution used to define regions in Ref.[3] is *very* closely related to the likelihood $\mathcal{L}(\rho)$ (thanks to Bayes' Rule). So, under most circumstances, regions assigned through a sensible interpretation of Ref.[3] will be quite similar to LR regions (albeit perhaps a bit larger). This suggest that further research may synthesize both approaches into a single, uniquely satisfying definition of "error bars".

ACKNOWLEDGEMENTS: This work was supported by the US Department of Energy through the LANL/LDRD program. I am grateful for helpful conversations over a long period of time with colleagues including (but not limited to) Chris Ferrie, Josh Coombs, Steve Flammia, Carl Caves, and Aephraim Steinberg. I am especially indebted to Matthias Christandl and Renato Renner for discussions about their closely related work, which inspired Lemma 1.

---

* Electronic address: robin@blumekohout.com

[1] Paris, M. and Řeháček, J. (eds) *Quantum state estimation* (Springer Verlag, 2004).

[2] Home, J. P. et al. Complete methods set for scalable ion trap quantum information processing *Science* **325**, 1227-30 (2009).

[3] Christandl, M. and Renner, R., Reliable quantum state tomography, arXiv:1108.5329,(2011).

[4] Feldman, G. J. and Cousins, R. D. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D* **57**, 3873–3889 (1998).

[5] Evans, S. N., Hansen, B. B., and Stark, P. B. Minimax expected measure confidence sets for restricted location parameters. *Bernoulli* **11**, 571–590 (2005).

[6] Robert, C.P. *The Bayesian choice: from decision-theoretic foundations to computational implementation* (Springer Verlag, 2007).

[7] Kempthorne, P. J., Numerical specification of discrete least favorable prior distributions. *SIAM J. Sci. Stat. Comp.* **8**, 171 (1987).

[8] Sugiyama, T., Turner, P. S., and Murao, M., Error probability analysis in quantum tomography: A tool for evaluating experiments. *Phys. Rev. A* **83**, 012105 (2011).

[9] Liu, Y.-K., Universal low-rank matrix recovery from Pauli measurements, arXiv:1103.2816,(2011).

[10] Kumar, P. and Yildirim, E. A., Minimum-volume enclosing ellipsoids and core sets. *J. Optimization Theory & Applications* **126**, 1-21 (2005).

[11] Fischer, K, Gartner, B, and Kutz, M. Fast smallest-enclosing-ball computation in high dimensions. *Lecture Notes in Computer Science* **2832**, 630-41 (2003).

[12] Christandl, M., König, R., and Renner, R., Postselection technique for quantum channels with applications to quantum cryptography. *Phys. Rev. Lett.* **102**, 020504 (2009).

[13] This is done by identifying a quantum process $\mathcal{E}$ with a bipartite state $\rho(\mathcal{E}) = (\mathbb{1} \otimes \mathcal{E})\,[|\Psi\rangle\langle\Psi|]$, where $|\Psi\rangle$ is a maximally entangled state. The differences between state and process tomography are entirely in implementation, and they all pertain to *gathering* data, not to its analysis.

[14] To see this, consider a simple biased point estimator that assigns $\hat{\rho} = \mathbb{1}/d$ no matter what data is observed. Clearly, the variance of $\hat{\rho}(D)$ is zero – and, just as clearly, this implies nothing about our uncertainty about the true ρ! While obviously extreme and even absurd, this estimator nonetheless demonstrates that the variance of biased estimators cannot be relied upon to describe uncertainty. In real-world tomography, the maximum likelihood estimator $\hat{\rho}_{\mathrm{MLE}}$ is biased by the positivity constraint $\rho \geq 0$, and this *can* lead to significant underestimation of uncertainty.

[15] This condition defines a Bayesian *credible region*.

[16] This bound was inspired by Refs.[3,12], and by discussions with Matthias Christandl, for whose help I am grateful!

[17] It's worth being precise here. Eq. 12 is not a lower bound on $F()$. Instead, the true exact $F()$ depends on ρ, and fluctuates above and below the χ^2 approximation, as shown in Fig. 4. Since $F()$ is greater than the χ^2 approximation for some ρ, any state-independent upper bound on $F()$ will be strictly greater than the χ^2 approximation, so Eq. 12 is a lower bound on such upper bounds.

[18] The authors of Ref.[3] and I became aware of each others' work only when our simultaneous submissions to the QIP 2012 conference were accepted and merged by the program committee. The present paper is *almost* independent of our subsequent collaboration; the exception is Lemma 1, which was proved by myself but directly inspired by conversations with Matthias Christandl at ETH-Zurich.